

Partitioned System-Defined Aggregate Tables

AtScale Cube Architects provide prioritized partitioning hints to the AtScale Engine to inform the creation of partitioned aggregate tables. The actual partitioning scheme used by the Engine depends on a number of factors, including:

- ▲ Whether the aggregate includes a column that matches a partition hint.
- ▲ Whether AtScale statistics suggest that partitioning would be worthwhile.
- ▲ Whether the target data warehouse supports table partitioning.

Procedure

1. In the Design Center, display the cube you want to work with.
2. In the right-hand tool panel, click the Aggregate Partition Scheme icon, which resembles a percent sign.
3. In the Aggregate Partition Scheme dialog box, click the menu icon on the right to add a partition.
4. In the Add Partition dialog box, select one or more columns to add. Click the folder for a dimension's attributes to display and select the columns, or use **Expand All** and **Collapse All** to work with columns.
5. Click **Save Partition**. Then add more partitions for other dimensions if you like.
6. Keep the **Include Higher Levels** checkbox selected. Including the higher levels of the partition can increase the usefulness of the partition at little cost.
7. Mouse over and click the left-hand dots of a partition to move it up or down the list.
8. Click the ellipses to the right of an added partition to remove it. You are prompted to confirm the deletion.
9. When you are done setting up partitions, click Publish at the far right to publish your changes.



Performance Best Practice: MDX Tools like Excel will specify filters using Attribute Key bindings. However, SQL and DAX-based tools like Tableau will specify filters using Attribute Name (Value) bindings. Data architects should be aware of their user community's tool usage so they may adopt a partitioning strategy that maximizes their query performance. If a cube services a mix of BI Tools, then it is recommended that Data Architects either partition by both name and key columns or use the same physical column for the AtScale Key and Name attribute bindings.

What To Do Next

Once a query containing the dimension levels that you've specified as a partition has been executed on the published cube, the AtScale engine determines whether to build a new system aggregate with the partition, based on the following:

- ▲ If the estimated number of rows per partition meets or exceeds the threshold set in the engine setting `AGGREGATES.CREATE.PARTITION.SYSTEMDEFINEDAGGREGATE.THRESHOLD`.
- ▲ If the estimated number of partitions in the aggregate table exceeds the threshold set in

`tables.create.partitions.maximumEstimatedNumberOfPartitions` the table will not be partitioned. The default value of this setting is 800.

Navigate to **Aggregates > Definitions** for the data warehouse to see whether the system aggregate was created as a result of your query. You can specify a level of a dimensional hierarchy to use as a partition key for partitioning instances of system-defined aggregate tables.